

ONEVM Accounting and Gratia (status report)

Tanya Levshina

Parag Mhashilkar

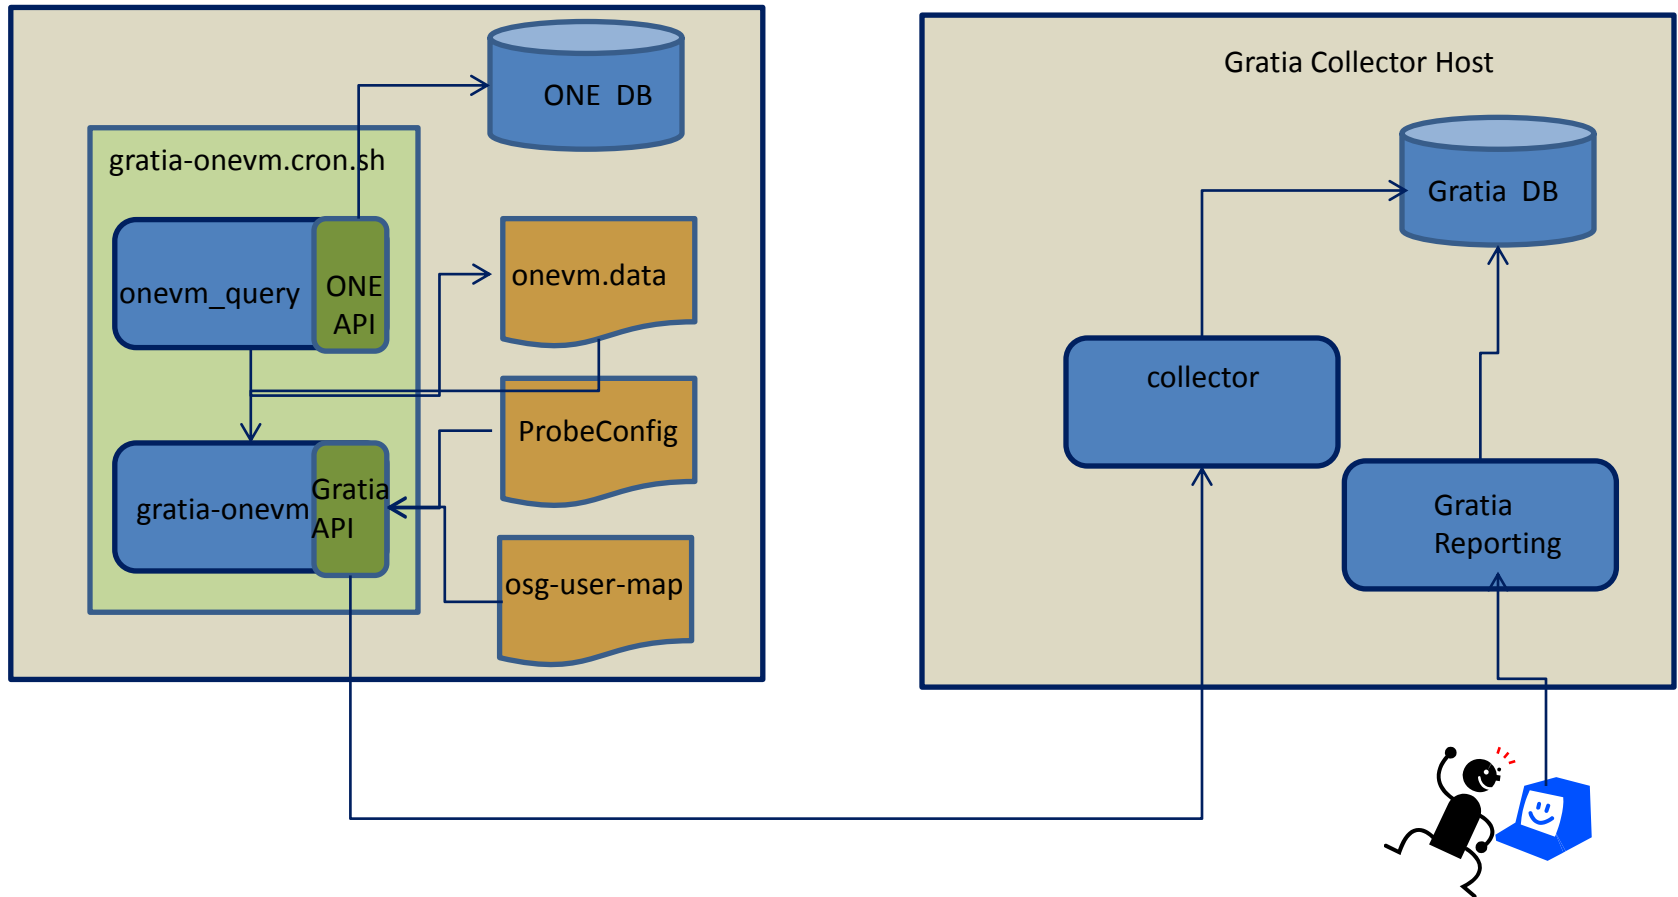
Requirements

- Provide accounting for Virtual Machines running on FermiCloud
- Use Gratia framework to collect usage records and extend the framework as required

Initial Assumptions

- All information about VMs prior and current states, start and stop times, resource usage etc could be acquired from interfaces provided by a VM management toolkit
- OGF Usage Record structure is currently sufficient to represent VM accounting information
- No changes are required to core Gratia Service at least during the initial phase
- The first accounting VM probe will be developed for Open Nebula 3.0.

Proposed Architecture



Developed Tools

- Tool (`one_query.rb`) gathers usage information from ONE using ONE API.
- Tool (`VMProbe.py`) reads this data and sends Usage Record to Gratia Collector using Gratia API.
- Script (`gratia-probe-onevm.cron.s`) orchestrates the execution of these tools. It runs as a cronjob.

ONE Data Mapping to OGF UR

ONR VM	OGF UR
NAME	JobName
ID	LocalJobId
Username	LocalUserId
NAME-ID#Instance#timestamp	GlobalJobId
MEMORY_REQ(MB)	Memory
VCPU	Processors
HOSTNAME	SubmitHost
IP	MachineName
HISTORY_STIME	StartTime
HISTORY_ETIME or current time	EndTime
HISTORY_ETIME-HISTORY_STIME	WallDuration, CpuDuration (user)

ONE Data Mapping to OGF UR (II)

ONE VM	OGF UR
STATE_STR:Failed (STATE_STR:DONE,HISTORY_REASON:None) (STATE_STR:DONE,HISTORY_REASON:1) (STATE_STR:DONE,HISTORY_REASON:0) (STATE_STR:DONE,HISTORY_REASON:2-4) (STATE_STR:Active,HISTORY_REASON:2-4) (STATE_STR:Active,HISTORY_REASON:0)	Status, Description 1 "VM exit code" 0 "VM exit code, success" 0 "VM exit code, migrated due to <> request" 0 "VM exit code, migrated due to <> request " 0 "VM still running"
"ONEVM"	ResourceType

Gratia Collector

- No changes in Gratia Collector were required
- JobUsageRecord, JobUsageRecord_Meta are populated

```
mysql> select dbid,LocalJobId, LocalUserId, CommonName, JobName, Status, StatusDescription, CPUUserDuration,WallDuration,
StartTime,EndTime,MachineName,SubmitHost from JobUsageRecord where ResourceType='ONEVM' and LocalUserId='parag' and
CommonName like '%cms%' and EndTime>='2011-12-12' order by dbid;
```

dbid	LocalJobId	LocalUserId	CommonName	JobName	Status	StatusDescription	CPUUserDuration	WallDuration	StartTime	EndTime	MachineName	SubmitHost
64796	195	parag	Generic cms user	test-3	0	VM is still running	3599.84	3599.84	2011-12-11 23:03:15	2011-12-12 00:03:15	192.168.154.163	fcl321
64806	192	parag	Generic cms user	test-3	0	VM is still running	3599.84	3599.84	2011-12-11 23:03:15	2011-12-12 00:03:15	192.168.154.160	fcl323
64809	198	parag	Generic cms user	test-3	0	VM is still running	3599.84	3599.84	2011-12-11 23:03:15	2011-12-12 00:03:15	192.168.154.155	fcl323


- Existing triggers and stored procedure populate summary table

```
select * from MasterSummaryData where ResourceType='ONEVM' and EndTime>='2011-12-12' order by SummaryID;
```

SummaryID	EndTime	VOcorrid	ProbeName	CommonName	DistinguishedName	ResourceType	HostDescription	ApplicationExitCode	Njobs	WallDuration	CpuUserDuration	CpuSystemDuration	Grid	Cores
2300	2011-12-12 00:00:00	52	onevm:fcll319.fnal.gov	Generic fermilab user		ONEVM			0					
520	1870132.789999999	1870132.789999999		0	OSG-ITB	1								
2301	2011-12-12 00:00:00	62	onevm:fcl319.fnal.gov	Generic cms user		ONEVM			0					
80	288000.32	288000.32		0	OSG-ITB	1								

Report Generation (I)

Minimal changes were needed to get reports for ONEVM probe (but these changes will only work for dedicated collector)




Report Generation with Webtools

Webtools from OSG repo installed on fermicloud025fnal.gov. No changes needed just configuration and access rights to mysql db.

[Gratia Events Graphs](#) [Data Queries](#) [Pie Chart](#) [Glidein ear Graphs](#) [BS](#)

Wall Hours by VO (Sum: 9,756 Hours)
14 Days from 2011-11-28 to 2011-12-12




Unknown (959.00)

cms (576.00)

Value
-12-12 21:26:16

Job Count by VO
14 Days from 2011-11-28 to 2011-12-12



fermilab

cms

Unknown

Maximum: 775.00 , Minimum: 0.00 , Average: 198.00 , Current: 660.00

Packaging and Deployment

- Work is done on gratia development branch
- gratia.spec for probes is modified so gratia-probe-onevm rpm can be built together with all other probes
- Probe is running on fcl319 for more than a week (with 1 hour frequency)

Known Problems

- ONE API doesn't provide means for filtering, so right now we are dumping ALL the data from ONE database. It takes about 1.2 minute for the time being. We have the following choices:
 - Put the request to Open Nebula developers and wait
 - Write our own API
 - Use ONE DB directly
- DN information is not available via API and it is not clear if it will be possible to extract DN from DB after the next release. It means that we can not specify Common Name in Gratia UR.
- Work is needed if we want report to be generated via BIRT when using common collector.

Next Steps

- Collaboration with ONE developers:
 - Establish working contact with ONE developers
 - Solve the filtering problems for ONE interface (effort to extend interface ~ 1 month)
 - Get information about DN, this should be available via interface
- Modify BIRT for reporting so it could be available from non-dedicated collector
- Provide more generic probe and plugin abilities so the probe would be suitable for any type of cloud management service.
- ONE API should be packaged as rpm